IBM Technical Brief

SAP® HANA® DB Migration from x86 to POWER® via HANA System Replication

Mark Gordon

IBM POWER and SAP Solutions

Version: 1.0
Date: February 28, 2018

1. Disclaimers

This paper demonstrates a way to migrate a HANA® database from x86 HANA 2.0 SPS02 to POWER HANA 2.0 SPS02 using HANA System Replication while minimizing the impact on end users. It is not a best practices guide. This process has been tested on our lab systems in the configuration described.

2. Trademarks

SUSE ® is a registered trademark of SUSE LLC.
SAP ®, SAP NetWeaver ®, and SAP HANA ® are registered trademarks of SAP SE.
Intel® is a registered trademark of Intel Corporation.
IBM® and POWER® are registered trademarks of IBM Corporation.

3. Version Changes

Version 1.0: February 28, 2018 – initial version

4. Acknowledgements

Thank you to Lou Lamprinakos, Damir Rubic and Andrew Castillo for reading the draft and offering suggestions for improvements.

5. Feedback

Please send comments or suggestions for changes to gordonmr@us.ibm.com.

6. Introduction

This paper demonstrates a process to migrate a scale-up HANA 2.0 DB on Intel to a HANA DB on IBM POWER using HANA System Replication. *The only end-user impact is a few minutes suspension of the SAP system when the takeover to HANA DB on IBM POWER is executed.*

HANA system replication can be used to migrate between X86 and POWER when the source DB is HANA 2.0, and the target DB is at an equal or later version of HANA.

HANA 2.0 on POWER runs on the Little-Endian (LE) SLES12 Operating System as well as on Red Hat RHEL for SAP HANA 7 (LE). HANA on X86 is also LE. *As described in SAP note 1999880 item 14, replication can be done between different types of hardware, as long as they are both LE.* In this demonstration, the source DB is on SLES 12 on X86 and the target is SLES 12 on POWER.

When using HANA System Replication, the primary and secondary (i.e. source and target) DBs must have similar topology - both must also have the same number of index servers.

SAP note 1984882 outlines how to use HANA system replication to change hardware platforms. The process is executed with more detail in this paper.

SAP note 2364819 describes a process to quiesce the application server work processes, so that one can do a near-zero downtime cutover from x86 HANA to POWER HANA. Do not use the quiesce process in SAP note 1913302, it has been superseded by 2364819.

We will apply SAP note 2381360, which supplies the program SUSPEND_HDB_CONNECTION used in SAP note 2364819. This program is used to suspend the application server work processes and disconnect them from the HANA database.

7. DB Migration Steps

Our source system is a single-node X86 (under VMware) HANA DB. It is not currently configured with HANA System Replication. It contains the HN1 tenant, which is a Netweaver 7.5 DB.

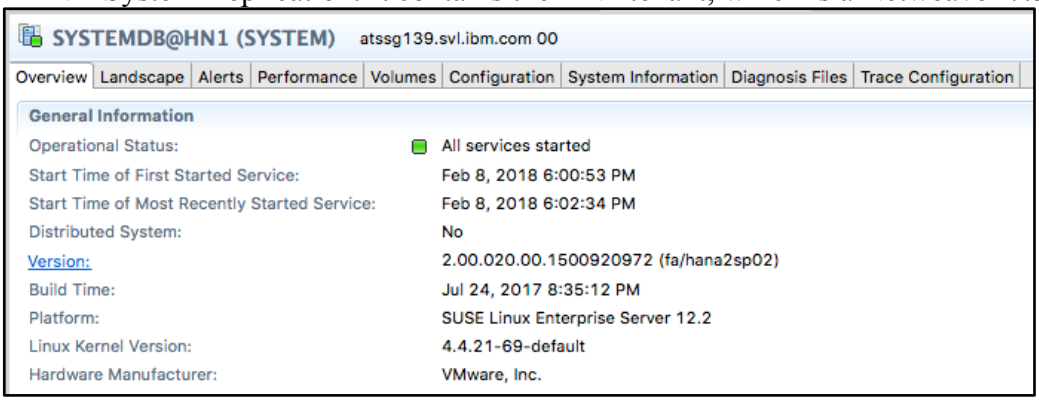

Figure 1: Current X86 productive HANA DB

Our goal is to migrate the contents of the X86 HANA DB to this POWER HANA DB system. This DB is a newly installed HANA DB.

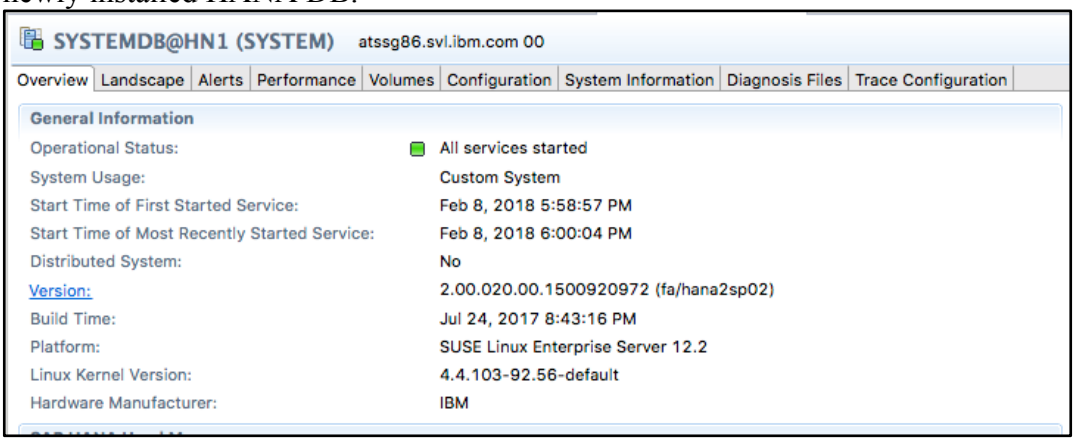

Figure 2: To-be productive POWER HANA DB

Check SAP note 1999880 item 15. The secondary must be at the same or higher patch level. In our case, both systems are the same version.

7.1. Pre-requisites

Verify support for the SAP application on HANA on POWER DB in the SAP PAM and SAP notes 2218464 and 2420699. Review OS support for HANA in SAP note 2235581.

7.2. Preparation

Apply SAP note 2381360 and transport it through the landscape to the Netweaver system that is being migrated.

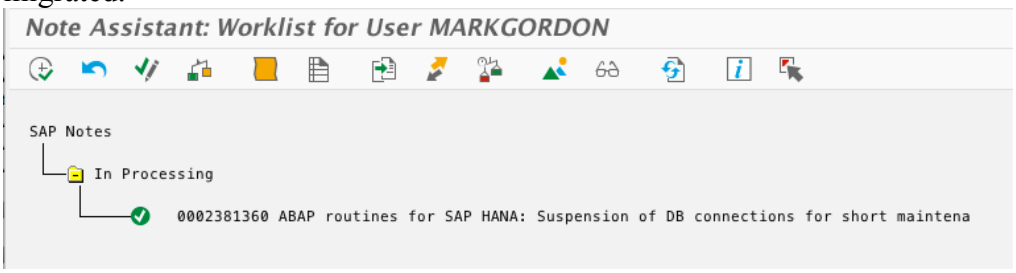

Figure 3: SNOTE

On the Netweaver SAP system, set SAP dbs/hdb/quiesce_all_connections parameters as described in SAP note 2364819. Also, rdisp/wpmax_run_time must be set to a value high enough to complete all the cutover activities in Section 7.4. We placed both in DEFAULT.PFL.

```
#
dbs/hdb/quiesce_all_connections = 1
rdisp/max_wprun_time = 600
```

Figure 4: Set SAP parameters for cutover

7.3. Configure System Replication

Our starting point is that host atssg139 is a running database for Netweaver 7.5, and host atssg86 is a newly installed HANA DB.

Per SAP note 2369981, we copy “key” and “dat” files to the secondary (POWER DB) system.

```
[atssg86:/usr/sap/HN1/SYS/global/security/rsecssfs/key # scp atssg139:/usr/sap/HN1/SYS/global/security/rsecssfs/key/SSFS_HN1.KEY .
[Password:
 SSFS_HN1.KEY
```

Figure 5: Copy KEY

```
[atssg86:/usr/sap/HN1/SYS/global/security/rsecssfs/data # scp atssg139:/usr/sap/HN1/SYS/global/security/rsecssfs/data/SSFS_HN1.DAT .
[Password:
 SSFS_HN1.DAT
```

Figure 6: Copy DAT

Enable the X86 DB system atssg139 as the primary (source) replication DB.

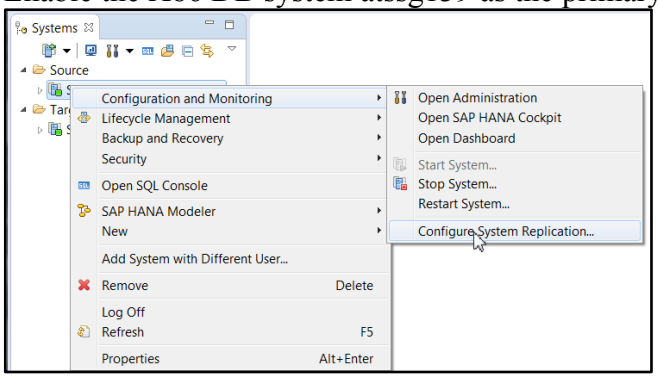

Figure 7: Configure replication on primary

Specify this will be the primary.

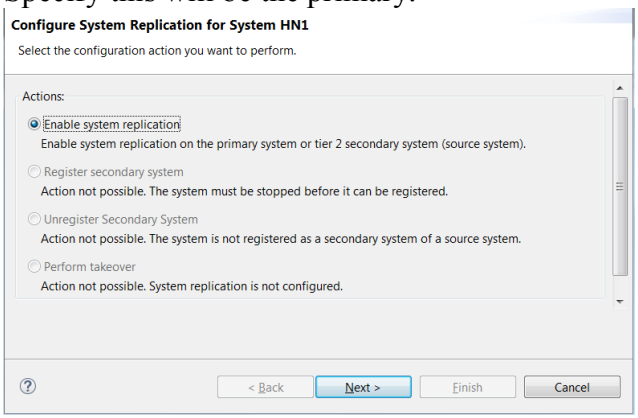

Figure 8: Set Primary

For replication, each system has a unique logical name. We will use “ASIDE” on atssg139 (x86) and “BSIDE” on atssg86 (POWER).

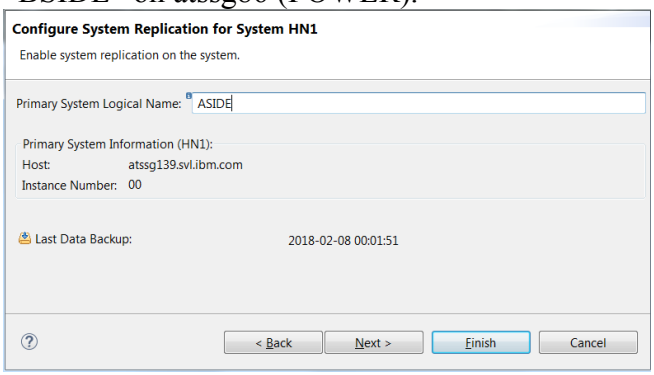

Figure 9: Parameters for Primary

After this executes, atssg139 is enabled to act as primary, but it has no secondary system attached to it.

Configure atssg86 as a secondary system. All the data in atssg139 will be replicated to atssg86. Shut down HANA on the to-be secondary before configuring replication.

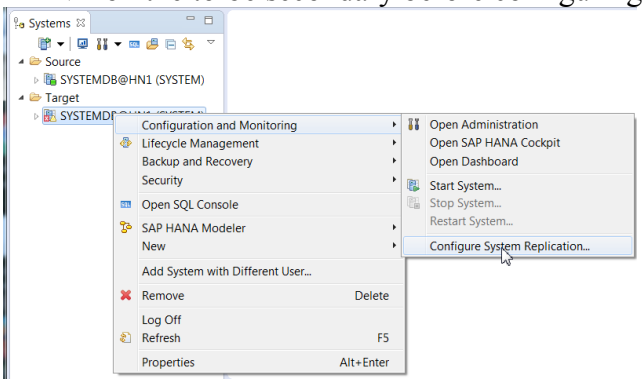

Figure 10: Configure replication on secondary

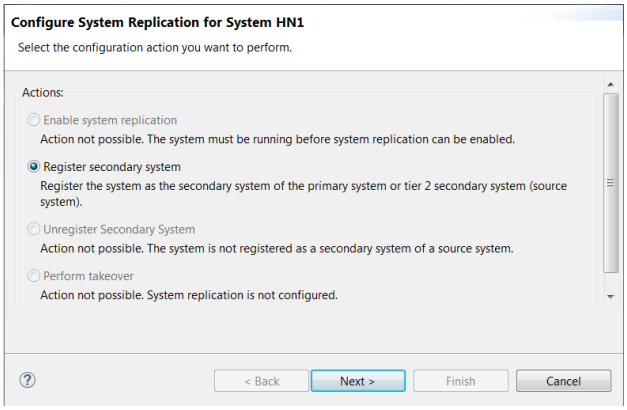

Figure 11: Register secondary

Enter parameters for registration.

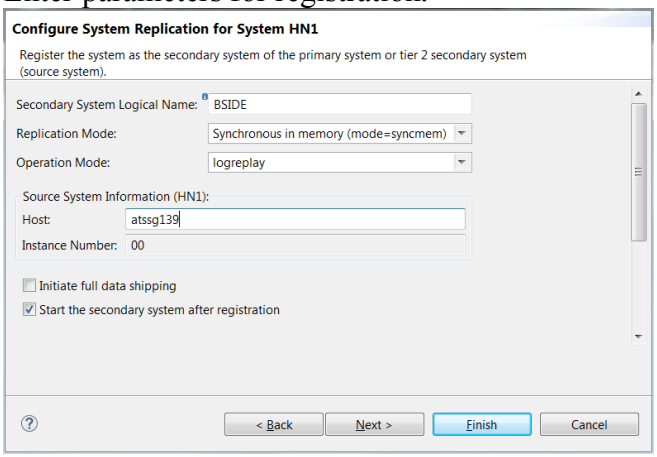

Figure 12: Secondary registration parameters

Wait for replication to be complete. The larger the database, the longer this will take.

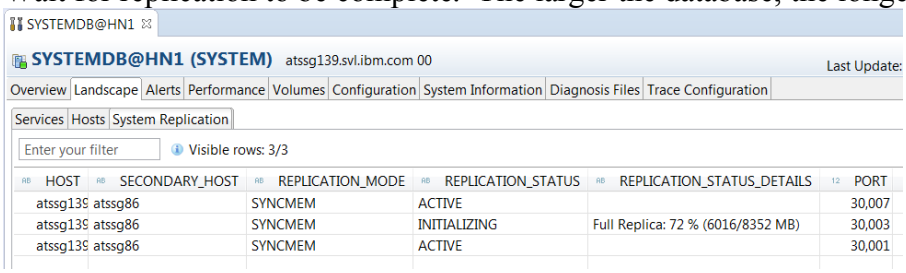

Figure 13: Waiting for replication to be caught up

Replication is complete.

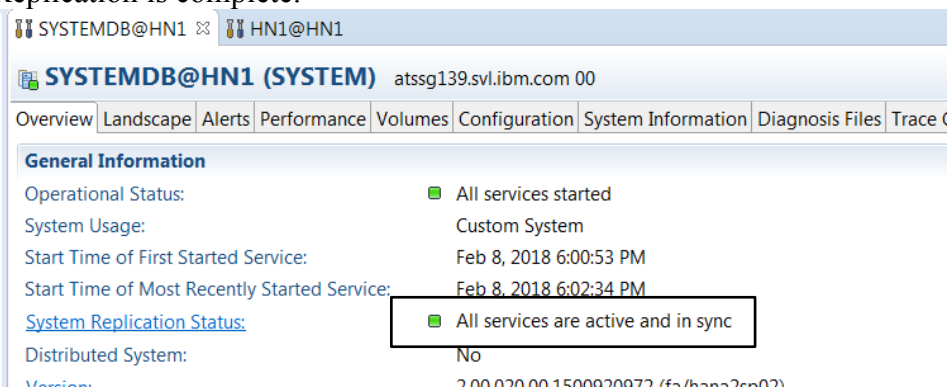

Figure 14: Replication active and in sync

7.4. Cutover to use POWER DB as productive system

Each SAP work process has at least one connection to the HANA DB. We will use the method in SAP note 2364819, and the program SUSPEND_HDB_CONNECTION to stop the processes during upgrade. Note that SUSPEND_HDB_CONNECTION suspends work process as they finish a dialog step. Long running batch jobs should *not* be scheduled to run at the time of the cutover to POWER HANA.

Using the SAP schema userid (SAPABAP1 on our test system), show the session list. The “ABAP:BED” Applications are SAP work process connections.

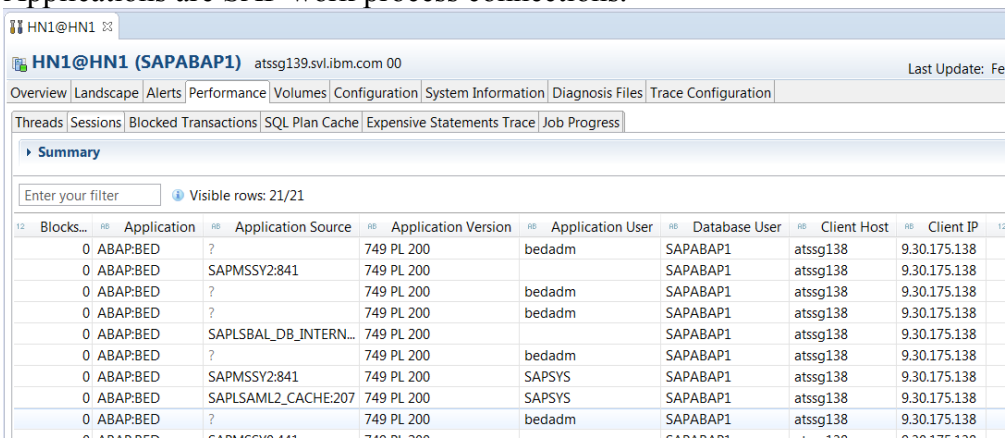

Figure 15: Session list

On the application server, create the file /usr/sap/<SID>/SYS/global/hdb_quiesce.dat. When the dbs/hdb/quiesce_check_enable parameter is set by SUSPEND_HDB_CONNECTION in the next step, the work processes will be stopped.

```
[atssg138:bedadm 16> touch /usr/sap/BED/SYS/global/hdb_quiesce.dat
 atssg138:bedadm 17>
```

Figure 16: touch hdb_quiesce.dat

On the Netweaver system, execute SUSPEND_HDB_CONNECTION.

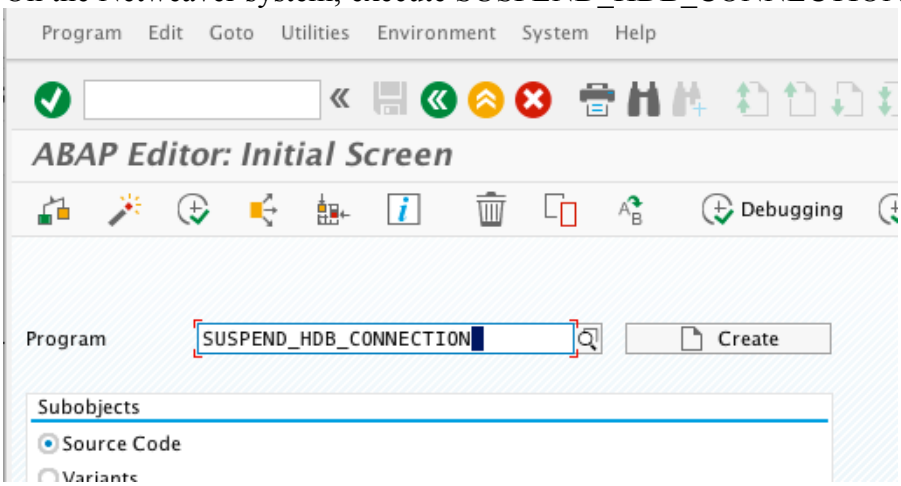

Figure 17: SUSPEND_HDB_CONNECTION

Set the parameters to suspend all work processes on all app servers.

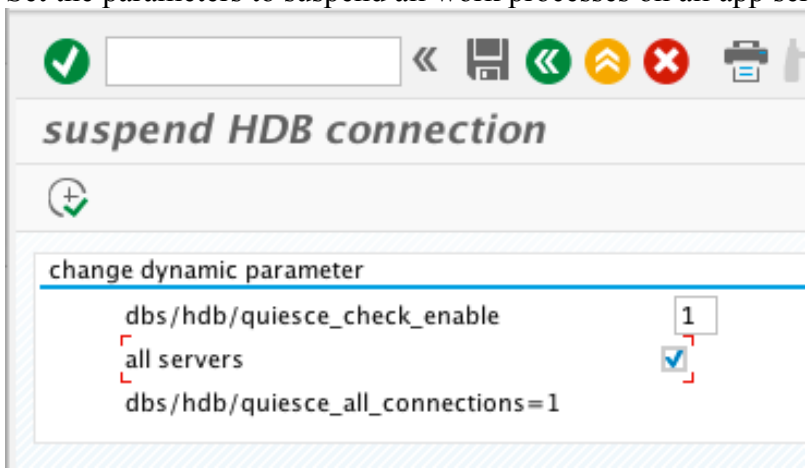

Figure 18: SUSPEND_HDB_CONNECTION parameters

When the program is executed, it hangs – all the work processes are suspended! We will need to check in HANA if there are other ABAP sessions.

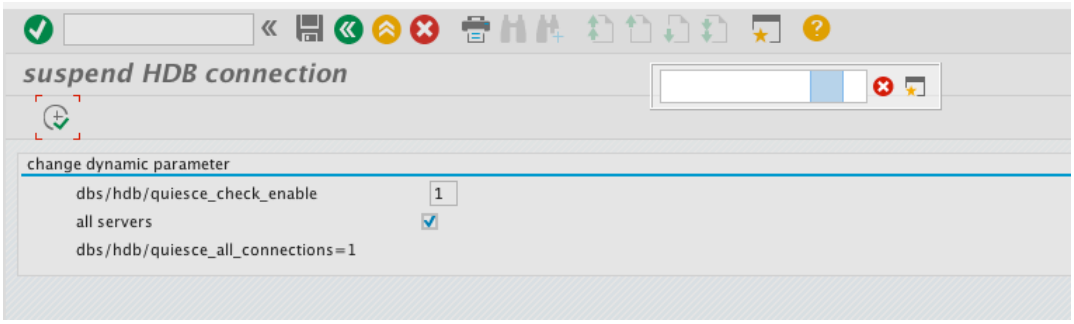

Figure 19: SUSPEND_HDB_CONNECTIONS hangs

Check using the session list in HANA. Note that all the "ABAP:BED" Applications are gone.

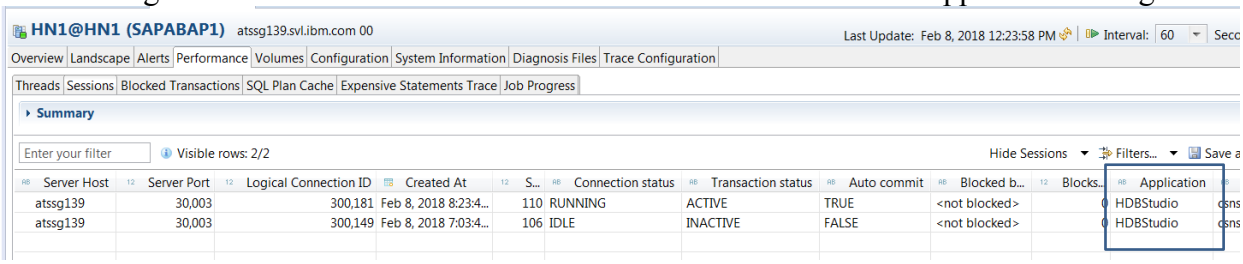

Figure 20: No ABAP sessions

We are now ready to make the POWER DB the productive database.

A virtual IP address, atssg137, is used by the application servers to identify the HANA DB.

```
[atssg138:bedadm 44> hdbuserstore LIST
DATA FILE       : /home/bedadm/.hdb/atssg138/SSFS_HDB.DAT
KEY FILE        : /home/bedadm/.hdb/atssg138/SSFS_HDB.KEY

KEY DEFAULT
  ENV : atssg137.svl.ibm.com:30015
  USER: SAPABAP1
```

Figure 21: hdbuserstore LIST on application server

Display network interfaces on atssg139, the X86 HANA DB. Note that atssg137 is a virtual IP address defined on adapter eth0.

```
atssg139:~ # netstat -i
Kernel Interface table
Iface   MTU Met    RX-OK RX-ERR RX-DRP RX-OVR    TX-OK TX-ERR TX-DRP TX-OVR Flg
eth0           1500   0  3030382      0      0      0  8557370      0      0      0 BMRU
eth0:atssg137  1500   0     - no statistics available -                        BMRU
lo            65536   0  5975827      0      0      0  5975827      0      0      0 LRU
atssg139:~ #
```

Figure 22: atssg139 netstat before remove VIPA

In the SUSE administration tool, YaST, delete the VIPA from atssg139.

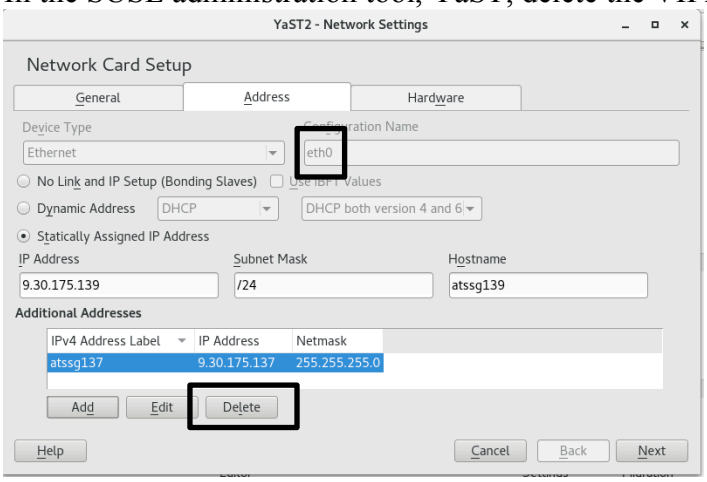

Figure 23: Delete VIPA on atssg139

Using command line on atssg139, verify VIPA atssg137 is gone. Once the VIPA is gone, it is not possible for the SAP work processes to connect to atssg137, until we re-create the VIPA after takeover.

```
atssg139:~ # netstat -i
Kernel Interface table
Iface   MTU Met    RX-OK RX-ERR RX-DRP RX-OVR    TX-OK TX-ERR TX-DRP TX-OVR Flg
eth0   1500   0  3092768      0      0      0  8650033      0      0      0 BMRU
lo    65536   0  6143206      0      0      0  6143206      0      0      0 LRU
atssg139:~ #
```

Figure 24: atssg139 after delete VIPA

On the current secondary system (atssg86 POWER HANA), perform HANA takeover.

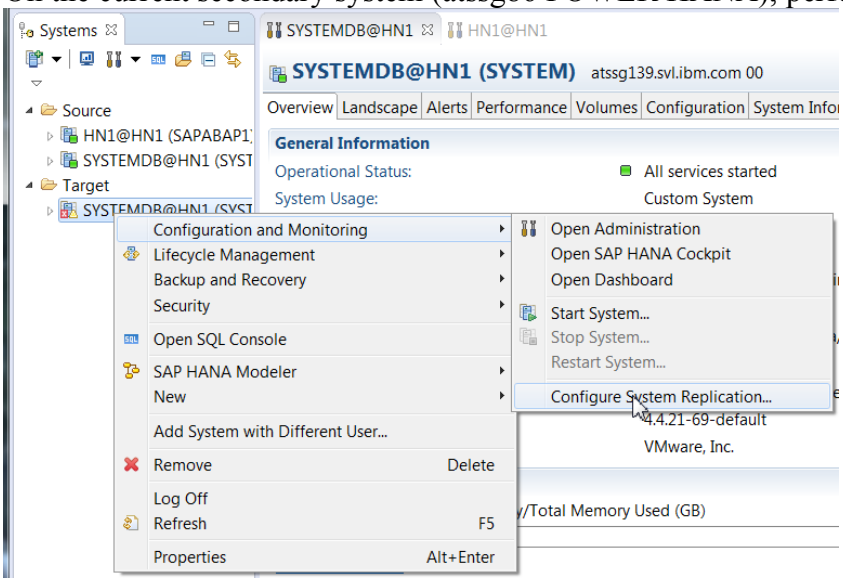

Figure 25: Start takeover on secondary

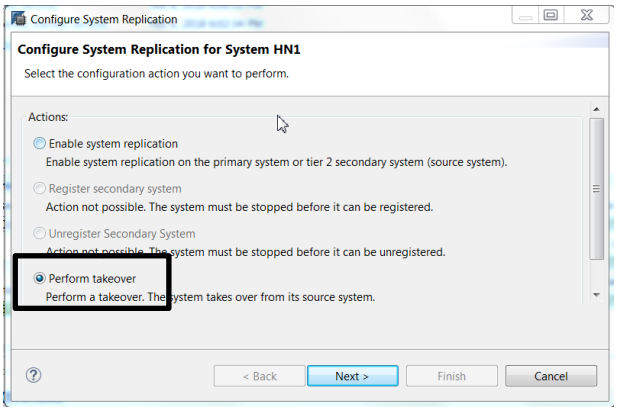

Figure 26: Perform takeover

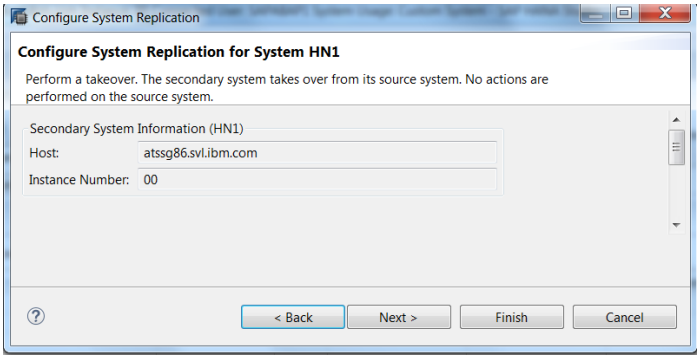

Figure 27: Takeover information

At this point, atssg139 (the former primary) is still running, but no SAP processes are connected. Also, atssg86 (the new productive DB) is up and ready to accept connections from the application server. But since the VIPA has been deleted and the work processes stopped, there are no connections to the atssg86 DB.

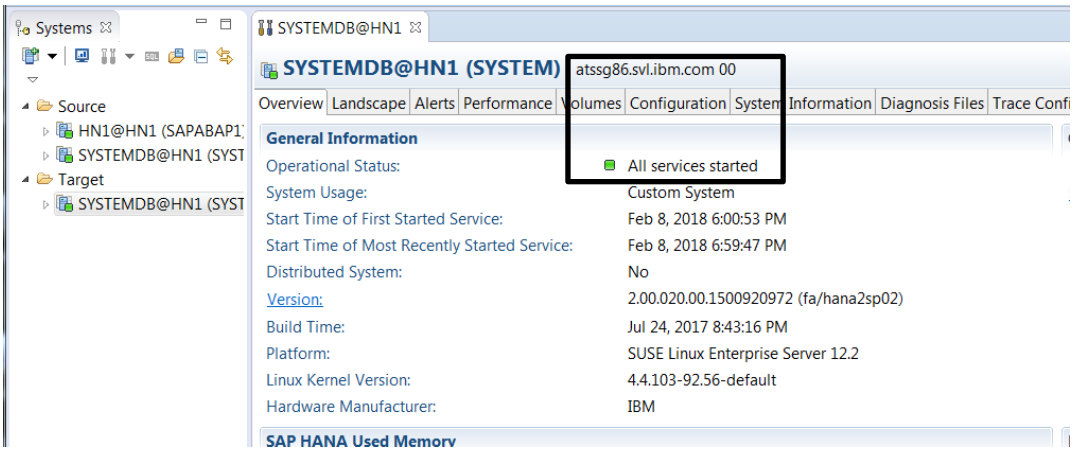

Shutdown the X86 DB server, to make sure no work processes can connect to the wrong system.

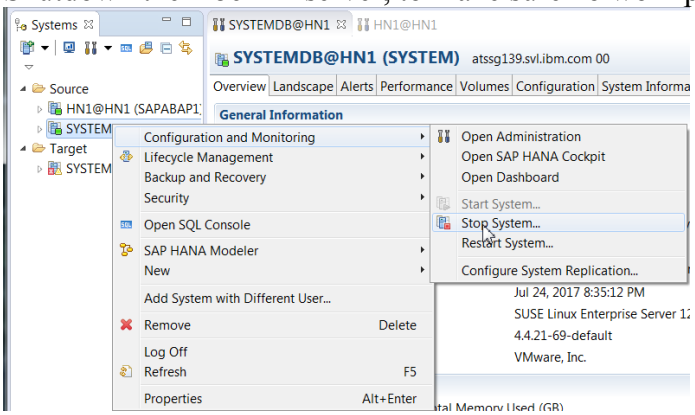

Figure 28: Shutdown HANA on atssg139

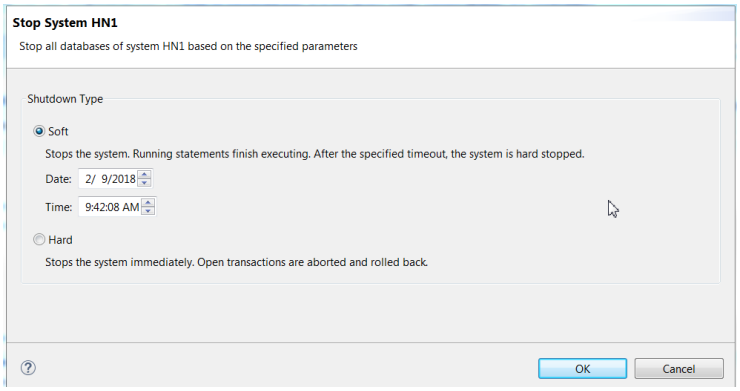

Figure 29: Stop System on atssg139

Once it is stopped, we go on to the next step.

Using SUSE YaST, define VIPA atssg137 on atssg86, so that the SAP work processes will be able to connect to the new DB server.

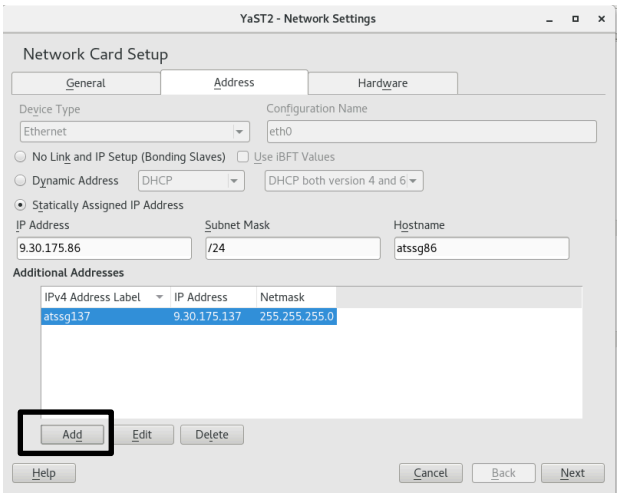

Figure 30: Add VIPA to POWER DB server

Verify atssg137 VIPA is now on atssg86.

```
atssg86:~ # netstat -i
Kernel Interface table
Iface   MTU Met    RX-OK RX-ERR RX-DRP RX-OVR    TX-OK TX-ERR TX-DRP TX-OVR Flg
eth0           1500   0  8298968      0      0      0  2306137      0      0      0 BMRU
eth0:atssg137  1500   0     - no statistics available -                        BMRU
lo            65536   0  4357048      0      0      0  4357048      0      0      0 LRU
atssg86:~ #
```

Figure 31: Confirm VIPA is on atssg86

Remove the file to enable work processes to restart

```
[atssg138:bedadm 51> rm /usr/sap/BED/SYS/global/hdb_quiesce.dat
```

The SUSPEND_HDB_CONNECTIONS program now finishes running

suspend HDB connection

```
suspend HDB connection

set  dbs/hdb/quiesce_check_enable  =  1

application server =  atssg138_BED_00                          : done

reconnect done
```

On the new productive DB, atssg86, show sessions for SAPABAP1. Note that the application server work processes ("ABAP:BED" Applications)have all reconnected.

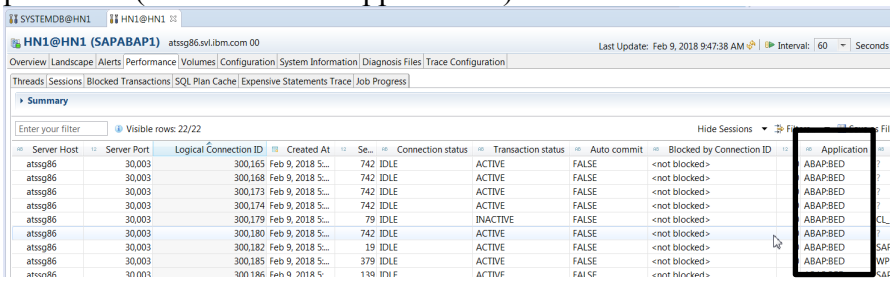

Figure 32: Session list on atssg86

We can now run transactions on the SAP application servers. DBACOCKPIT shows SAP is now connected to atssg86, the POWER DB server. We have successfully executed the migration from X86 HANA to POWER HANA.

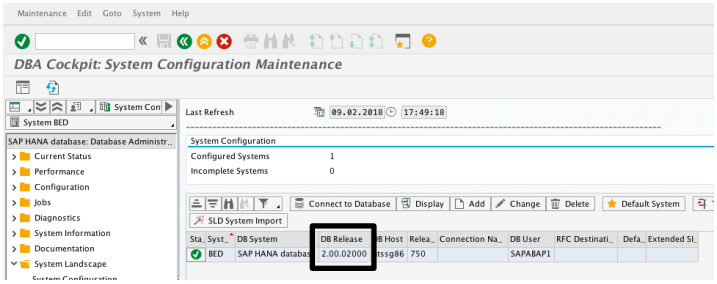

Run SUSPEND_HDB_CONNECTION and set dbs/hdb/quiesce_check_enable back to original value of 0.

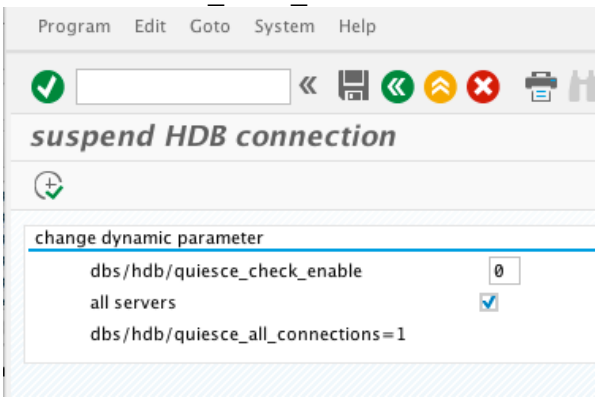

And at the next service window, remove "dbs/hdb/quiesce_all_connections" from the SAP parameter files.